Morgan Guaranty Trust Company
Risk Management Services
Peter Zangari (1-212) 648-8641
zangari_peter@jpmorgan.com

A VaR methodology for portfolios that include options

Risk managers who implement Value at Risk (VaR) systems frequently face the daunting task of measuring the risk of a portfolio that contains options. The nature of this problem results from the standard VaR assumption that portfolio return distributions are conditionally normal. Among other things, this implies that return distributions are symmetric. However, due to the payoff structure of options, many portfolios that include options have return distributions that are, at the very least, skewed.[1] In this article we suggest a modification to standard VaR computations that offer practitioners a means of estimating the risk of a portfolio that includes options.

In its standard context VaR estimates are given by the bands of a symmetric confidence interval around the expected value of a portfolio's return. These bands represent the largest expected change in the value of the portfolio with a specified level of probability. For example, if R_p is the return on a portfolio with mean $E(R_p)$, its 90% confidence interval is given by:

$$C_v = \{-1.65\sigma_P + E(R_P), E(R_P) + 1.65\sigma_P\}$$

where -/+ 1.65 are the 5th/95th percentiles of the standardized normal distribution. Over short horizons, the estimate of $E(R_p)$ is often set to zero to reduce the noise in estimating the sample mean.[2]

In general, when a portfolio's payoff is a nonlinear function of some underlying returns, even if these returns are distributed normally, the confidence interval estimate for the expected value of the portfolio using C_V is inappropriate. For example, portfolios with nonlinear payoffs may have skewed return distributions. Skewness invalidates the application of symmetry imposed by the scale factors +/- 1.65 (the quantiles of the standard normal distribution). In addition, nonlinearities transform the moments (e.g., mean, variance, skewness, etc.) of the underlying return distribution. Therefore, assumptions placed on the expected values of underlying returns do not necessarily carry over to a portfolio's expected values.

In order to properly evaluate the risk of a portfolio that contains nonlinear instruments, researchers often propose full simulation routines. According to this methodology, a path of future underlying prices are generated and the portfolio's value, which consists of options, is revalued at various prices along the path. A specific type of full simulation, known as Structured Monte Carlo, is outlined in the *RiskMetrics™ Technical Document*. A major disadvantage of the full simulation approach is that it is computationally and time intensive.

This study presents two methods to compute the Value at Risk estimates of portfolios with nonlinear payoffs that do not require full simulation. Its goal is to present a methodology that is relatively simple to implement and does not require a lot of computer time. The methodology is developed from first principles and is used to compute the VaR over a five-week horizon of the following position.[3]

On April 18, 1995, a U.S. dollar based investor buys a USD1,000,000 nominal value 2-year French franc government bond (OAT Strip) at a yield of 7.147%. In order to hedge FX exposure, the investor buys a 5-week (which corresponds to the investment horizon) FRF/USD put

[1] In fact an option's return distribution is a mixture of discrete and continuous variables.
[2] See *RiskMetrics™ Technical Document* for details.
[3] This position originally appeared in the *RiskMetrics™ Technical Document*, 3rd edition.

option on a notional USD 870,994 at the money forward FX rate of 4.864. The current value of the option is FRF/USD 0.04616. Therefore, it costs USD 8,289 to hedge USD 870,994.

Market risk is often analyzed in terms of its delta (1st order) and gamma (2nd order) risk. Below, in the case where VaR incorporates gamma risk, we compare results obtained using the methods described herein to those given by full simulation which serves as our benchmark. In so doing we highlight important differences and similarities among the different methodologies. To facilitate the discussion we will use the following definitions and parameter settings.

PV = Amount of current position in USD (present value) = USD870,994	$P_{X(t+1)}$ = any future value of $P_{X(t)}$
R_B = 5-week return on 2-year OAT	Current spot rate FRF/USD=4.855
σ_B = 5-week forecast standard deviation of bond price = 0.7757% (yield standard deviation=5.83%)	1-month forward FRF/USD rate=4.684
R_X = 5-week return on FRF/USD exchange rate	Current 2-year OAT yield=7.147%
σ_X = 5-week forecast standard deviation of FX rate = 3.117%	Current 2-year USD yield=6.125%
δ = 0.532	$\Gamma = 3.15$
R_O = return on FX option	$\rho_{B,X}$ = correlation of OAT and FX rate = -.291

We analyze the risk of the synthetic portfolio as follows. First, the option is ignored and we compute the Value at Risk of the position that consists only of the purchased OAT. We then introduce the option to hedge foreign exchange risk on the OAT. We focus on two specific types of market risk associated with holding the option – delta risk and gamma risk. Initially we focus exclusively on delta risk since the standard VaR methodology is still applicable. Next, gamma risk is added. Now, the nonlinearity of the option's payoff dominates the portfolio's return distribution. Since standard VaR methodology is no longer appropriate, we use a normal analytical approximation known as the Cornish-Fisher expansion to find the percentiles of this portfolio's distribution. These percentiles are then used to estimate VaR. Finally, we conduct an experiment to determine how the normal approximation performs. The method used in this experiment, which we refer to as partial simulation, actually turns out to be an alternative technique that may be used to compute VaR in the presence of gamma risk.

The unhedged position

Suppose the investor buys the OAT but does not hedge its foreign exchange exposure. In this case the return on the portfolio, R_p, is simply the sum of the returns on the FRF/USD and OAT which is written as $R_p = R_X + R_B$. It's standard deviation is:

[1] $$\sigma_1 = \sqrt{\sigma_B^2 + \sigma_X^2 + 2 * \rho_{B,X} * \sigma_B * \sigma_X}$$

Assuming that R_X and R_B are normal, as in standard VaR, we know that R_p is also normal so the Value at Risk of holding the foreign bond is

[2] $$\text{VaR1} = \text{PV} * (1.65) * \sigma_1 = \text{USD } 42,907$$

Hedging FX exposure

In an effort to reduce foreign exchange market risk the investor buys a put option on the FRF/USD exchange rate. Having purchased the put option, the return on the portfolio that also consists of the OAT is now $R_P=R_B+R_X+R_O$. In order to compute the return on the portfolio, R_P, we need an expression for the return on the option, R_O. We value the FX option using the Garman-Kohlhagen formula. Specifically, for a given set of parameters denote the option's value by $V(P_{X(t)},K,\tau,\rho,\sigma_X)$ where:

$P_{X(t)}$ = FRF/USD spot rate at time t

K = the option's exercise price

τ = time to maturity in terms of a year

ρ = riskless rate of a security that matures when the option does

σ_X = 5-week forecast standard deviation of FX rate

A first step toward obtaining an expression for R_O is to approximate the future value of the option $V(P_{X(t+1)},K,\tau,r,\sigma_X)$ with a 2nd-order Taylor series expansion around the current values (spot rates), $P_{X(t)}$, K, τ, r, and σ_X. This yields[4]

[3] $$V(P_{X(t+1)},K,r,\tau,\sigma_X) \approx V_o(P_{X(t)},\mathrm{K},\mathrm{r},\tau,\sigma_X) + \delta * (P_{X(t+1)} - P_{X(t)}) + \frac{1}{2}\Gamma * (P_{X(t+1)} - P_{X(t)})^2$$

which can be written more succinctly as $dV = \delta * dP_X + \frac{1}{2}\Gamma * (dP_X)^2$ where the option's delta (δ) and gamma (Γ) are equal to 0.532 and 3.15, respectively.

Note dV, the absolute change in the value of the option, is in units of P_X. This follows from the fact that δ is unitless and Γ is in units of $1/P_X$. Since RiskMetrics™ currently provides the volatility of returns we write dV as a function of relative price changes

[4]
$$\mathrm{dV} = \mathrm{P}_{\mathrm{x(t)}} * \delta * \left(\frac{\mathrm{dP_x}}{\mathrm{P_{x(t)}}}\right) + \frac{1}{2} * \mathrm{P}^2_{\mathrm{x(t)}} * \Gamma * \left(\frac{\mathrm{dP_x}}{\mathrm{P_{x(t)}}}\right)^2$$
or
$$\mathrm{dV} = \mathrm{P}_{\mathrm{x(t)}} * \delta * (\mathrm{R_x}) + \frac{1}{2} * \mathrm{P}^2_{\mathrm{x(t)}} * \Gamma * (\mathrm{R}_x)^2$$

Here, dV relates an absolute change in the value of the option to a relative change in the foreign exchange rate. However, since dV is still in units of P_X we need to standardize it to make it unitless. This allows us to obtain the relative return on the option

[5] $$\mathrm{R_o} = \delta * (\mathrm{R_x}) + \frac{1}{2} * \mathrm{P_{x(t)}} * \Gamma * (\mathrm{R_x})^2$$

Therefore, the return on the portfolio of the hedged position $R_P=R_B+R_X+R_O$ is

[6] $$\mathrm{R_P} = R_B + R_X + \delta * \mathrm{R_X} + \frac{1}{2} * \mathrm{P_{X(t)}} * \Gamma * (\mathrm{R_X})^2$$

[4] In the following expression, delta (δ) is the first derivative of V with respect to P_X and gamma (Γ) is the second derivative of V with respect to P_X.

Morgan Guaranty Trust Company
Risk Management Services
Peter Zangari (1-212) 648-8641
zangari_peter@jpmorgan.com

Since the investor is purchasing a put to hedge foreign exchange risk, $\delta < 0$. For ease of exposition we will use $\delta^* = |\delta|$ in the following analysis. Notice from [6] that we can compute two types of Value at Risk estimates. The first incorporates the linear risk of the option. In this case, we use only the first three terms of [6]. Alternatively, we can capture the nonlinear features of the option's return distribution by also including the gamma effect (the last term).

If the investor only wants to account for the delta component of the option, the return on the portfolio is:

[7] $$R_P = R_B + (1-\delta^*)R_X$$

which has a standard error

[8] $$\sigma_2 = \sqrt{\sigma_B^2 + (1-\delta^*)^2\sigma_X^2 + 2*(1-\delta^*)\sigma_{B,X}^2}$$

Assuming that both R_B and R_x are normal implies that R_p is also normal. The Value at Risk of this hedged position accounting only for the delta risk of the option is

[9] $$\text{VaR2} = \text{PV}*(1.65)*\sigma_2 = \text{USD } 20{,}698$$

Intuitively, it can be seen from [8] that when $\delta^*=1$, the FX risk is completely hedged and all that is left is interest rate risk. On the other hand, $\delta^*=0$ implies there is no hedge and the position is as if the investor held a foreign bond (see [1]). For values of δ^* between 0 and 1, the VaR of holding a foreign bond and FX option will be lower than if no option was held.

As previously shown, the portfolio's return that accounts for both the delta and gamma effect of the option is

[10] $$R_P^G = R_B + (1-\delta^*)R_X + \frac{1}{2}*P_{X(t)}*\Gamma*\left(R_X\right)^2$$

A consequence of including the term Γ is that R_P^G's distribution becomes right skewed. To see how the option's delta and gamma components effect the portfolio's return distribution, chart 1 presents probability density functions (pdf) for two portfolio return series. One pdf is based only on the delta component (see [7]), the other is based on both the delta and gamma components (see [10]).

Chart 1
The effect of incorporating gamma risk on a portfolio's return distribution

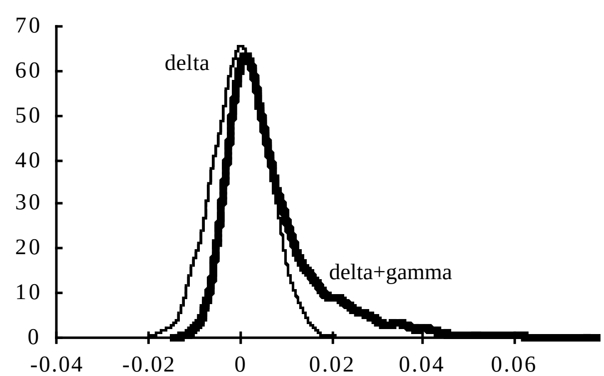

Morgan Guaranty Trust Company
Risk Management Services
Peter Zangari (1-212) 648-8641
zangari_peter@jpmorgan.com

A striking feature from chart 1 is the skewness embedded in the return distribution that includes the gamma effect. In fact, the distribution (grey line) that only accounts for delta is simply a scaled version of the normal distribution. Finally, for future use we need to derive the standard deviation of R_P^G which is presented below

[11] $$\sigma_3 = \sqrt{\sigma_B^2 + (1-\delta^*)^2 * \sigma_X^2 + 2*(1-\delta^*)*\sigma_{B,X}^2 + \frac{1}{2} * P_{X(t)}^2 * \Gamma * \sigma_X^4}$$

This expression[5] follows from the assumption that R_B and R_X are normal.

Accounting for the distributional features of R_P^G

In the presence of gamma risk, normal VaR methodology which relies on the critical values +/- 1.65 will give misleading risk estimates. The reason for this is simple: +/- 1.65 come from the normal distribution, however, as seen from chart 1, the gamma component of the option causes the portfolio return distribution to be highly skewed. The inaccuracy of normal VaR is shown in chart 2 below. We plot R_P^G's distribution (grey line) and that given by the normal distribution with a zero mean and a variance equal to σ_3^2 (black line).

Chart 2
A comparison of true VaR and standard VaR

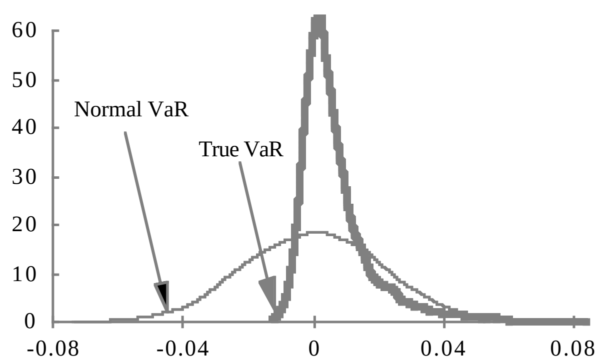

Essentially, by incorporating gamma risk the investor reduces his risk by the difference between true VaR and normal VaR. The discrepancy between true and normal VaR leads us to search for methods that augment the standard VaR methodology to account for the skewed return distribution. In particular, we seek counterparts to the quantiles +/- 1.65 that capture the skewness of R_P^G's distribution.

Focusing on analytical solutions, there are basically three approaches we could take to find the percentiles of R_P^G's distribution. First, we could match the moments of R_P^G to a general family of distributions (Pearson family), second, we could construct the distribution of R_P^G as a deformation of a standard normal variables, and, third, we could use the moments of R_P^G and a normal analytical approximation to estimate the percentiles of R_P^G. In this article we describe how to apply this last method. We find the critical points of R_P^G's distribution (i.e., the counterparts to -/+ 1.65) by applying a formula known as the Cornish-Fisher expansion.
Applications of normal analytical approximations are motivated by the understanding that any

any distribution can be viewed as a function of any other one. For example, the 5th and 95th percentiles of R_P^G's distribution denoted $cv_{.05}$ and $cv_{.95}$ can be calculated as a function of the standard normal percentiles $z_{.05}$ =-1.65 and $z_{.95}$ =1.65, and R_P^G's estimated moments. To be more specific, consider again the normal 90% confidence interval around the mean portfolio return $E[R_p]$

[12] $$C_v = \left\{-1.65\sigma_P + E(R_p), E(R_p) + 1.65\sigma_P\right\}$$

Under the maintained assumptions, when R_p is no longer normal, that is, when R_p becomes R_P^G we can write the approximate confidence interval for $E[R_P^G]$ as

[13] $$\begin{aligned} C_G &= \left\{E\left[R_P^G\right] + (-1.65 + s_{.05}) * \sigma, E\left[R_P^G\right] + (1.65 + s_{.95}) * \sigma\right\} \\ &= \left\{E\left[R_P^G\right] + (cv_{.05}) * \sigma, E\left[R_P^G\right] + (cv_{.95}) * \sigma\right\} \end{aligned}$$

The main purpose of the correction term s_α is to adjust for skewness. To a lesser extent it corrects for higher-order departures from normality. In the case of the normal approximation interval, $s_{.05}=s_{.95}=0$. In practice, the Cornish-Fisher expansion allows us to compute the adjusted critical values $cv_{.05}$ and $cv_{.95}$ as a function of the normal critical values $z_{.05}$ and $z_{.95}$ directly.[6]

[14] $$\text{cv} = z_\alpha + \frac{1}{6}(z_a^2 - 1) * \rho_3 + \frac{1}{24}(z_\alpha^3 - 3z_\alpha) * \rho_4 - \frac{1}{36}(2z_\alpha^3 - 5z_\alpha) * \rho_3^2$$

where

$\rho 3=E[(R_P^G-E[R_P^G])^3/\sigma^3$ measures R_P^G's skewness

$\rho 4=E[(R_P^G-E[R_P^G])^4/\sigma^4-3$ measures R_P^G's kurtosis

For example, if we wanted to compute the adjusted percentile $cv._{05}$ associated with -1.65, we would use:

[15]

$$\text{cv}_{.05} = -1.65 + \frac{1}{6}((-1.65)^2 - 1) * \rho_3 + \frac{1}{24}((-1.65)^3 - 3(-1.65)) * \rho_4 - \frac{1}{36}(2(-1.65)_3 - 5(-1.65))\rho_3^2$$

Under the assumption that returns are normally distributed, ρ_3 and ρ_4 can be written directly as a function of the variances and covariances of R_B and R_X. This result is very useful since multivariate extensions are straightforward and standard VaR calculations already require a covariance matrix. The measures ρ_3 and ρ_4 depend on the cumulants of R_P^G where the first four cumulants of R_P^G, denoted {κ1, κ2, κ3, κ4 } are defined as

[16] $$\begin{aligned} \kappa_1 &= E[R_P^G] \\ \kappa_2 &= \text{var}(R_P^G) = E[(R_P^G)^2] - E[(R_P^G)]^2 \\ \kappa_3 &= (E[R_P^G] - E[(R_P^G)])^3 \\ \kappa_4 &= (E[R_P^G] - E[(R_P^G)])^4 - 3\,\text{var}(R_P^G)^2 \end{aligned}$$

[6] In this article we present only the first 4 terms of the Cornish-Fisher expression. For a sample size n, this approximation has an error of order $O(n^{-3/2})$. For a more complete version, see Johnson and Kotz (1970).

Morgan Guaranty Trust Company
Risk Management Services
Peter Zangari (1-212) 648-8641
zangari_peter@jpmorgan.com

Using σ_B, σ_X, ρ_{BX}, δ, and Γ the cumulants of R_P^G are κ_1= 0.745%, κ_2= 0.0318%, κ_3=0.000844%, and κ_4= 0.00000109% and the standardized coefficients[7] are ρ_3=1.48 and ρ_4=0.107. Substituting these values into [14], R_P^G's lower ($cv_{.05}$) and upper ($cv_{.95}$) critical values are -1.176 and 2.029, respectively. Table 1 compares these values to those from the normal distribution. It is evident that the Cornish-Fisher approximation captures the skewness of R_P^G's distribution.

Table 1
Percentiles for normal and Cornish-Fisher approximation

Percentile	5th	95th
Normal	-1.650	1.650
Cornish-Fisher approximation	-1.176	2.029
Relative difference	+28.7%	+22.9%

Having calculated the adjusted percentiles, the 90% confidence interval for the expected return on the portfolio that consists of an OAT and a put option is:

$$C_G = \{-1.176\sigma_P + E(R_P^G), E(R_P^G) + 2.029\sigma_P\}$$

where $\sigma_p = 1.78\%$ and $E(R_P^G) = 0.745\%$. Using these results, the VaR of this portfolio is

[17] $$\text{VaR3} = PV * cv_{.05} * \sigma_3 = \text{USD } 18,285$$

When applying [15] it is important to remember that this expression is exact only when the true values of the standardized cumulants are used. In practice, however, we evaluate [14] using sample estimates of the standardized cumulants. When sample estimates are used to evaluate a mathematical expression we face what is known as a "certainty equivalence" problem. Essentially what happens is that the estimation error embedded in the sample estimates is carried over to the numerical value produced by [14]. Consequently, if there is a lot of estimation error, the Cornish-Fisher critical values (*cv*'s) will be inaccurate. To determine how estimation error affects the Cornish-Fisher approximation, we find the critical values of R_P^G by simulating its distribution and then finding the 5th and 95th percentiles. Note that this is not the same as full simulation mentioned earlier because here nothing is revalued. All that is required is that we generate a matrix of normal random numbers denoted Y and then apply [10].

This partial simulation approach works as follows. Let N denote the number of simulated random variables and define an Nx2 matrix of independent normal random variables Y=[Y1 Y2] where Y1 and Y2 are both N x 1 random. Using Y and the covariance matrix of R_B and R_X denoted Σ, simulate X=[R_B R_X], an Nx2 matrix of correlated normal random variables vectors.[8] Defining Δ=[1, 1-δ] (2x1), γ=[0,Γ] (2x1) and P_X as the spot FRF/USD exchange rate, the distribution of R_P^G is generated using the expression:

[18] $$R_P^G = X * \Delta + \frac{1}{2} P_x * X^2 * \gamma$$

[7] Exact formulae for the cumulants are provided in a *Technical Appendix* that is available from the author upon request
[8] See *RiskMetrics™ Technical document* (3rd edition) for details

Denoting the 5th percentile of the standardized distribution of R_P^G by $m_{.05}$, VaR under partial simulation is:

[19] $$\text{VaR4} = m_{.05} * \sigma_3 = \text{USD } 19{,}243$$

Table 2 summarizes the results of this section and presents VaR estimates produced by full simulation.

Table 2
VaR bands for various methodologies
USD

Confidence Interval 90%	Normal	Cornish Fisher	Partial simulation	Full simulation
Lower 5% (VaR)	-20,698	-18,285	-19,243	-19,596
Upper 5%	20,698	31,543	29,789	33,538

The Cornish-Fisher expansion and the partial simulation approach give similar results and both are an improvement over the normal model (in comparison to full simulation). In full simulation, the 5th percentile of the profit/loss distribution is USD -4,008. Since the mean of the distribution is USD 15,588, the adverse price move from the mean (VaR) is USD 19,596.

The expected value of a portfolio with gamma risk

In the previous section we established how accounting for the gamma effect $1/2 * P_{X(t)} * \Gamma * (R_X)^2$ skews the portfolio's return distribution. Another feature of including this term is that even if it is assumed that the underlying returns have a zero mean (i.e., $E[R_B]=E[R_X]=0$), the portfolio's expected value, $E[R_P^G]$, is not necessarily zero since

[20] $$E[R_P^G] = \frac{1}{2} * \Gamma * P_{X(t)} * \sigma_X^2 = 0.745\%$$

or $$E[R_P^G] = \text{USD } 6{,}488$$

Using this as the appropriate mean portfolio return, we can compare the lower and upper bounds of 90% confidence intervals generated by normal, Cornish-Fisher, partial and full simulation methods. These bounds are presented in table 3.

Table 3
VaR bands relative to expected portfolio return
Values correspond to VaR + mean portfolio return (in USD)

Confidence Interval 90%	Normal	Cornish Fisher	Partial simulation	Full simulation
Lower 5%	-14,210	-11,797	-12,755	-4,008
Upper 5%	27,186	38,031	36,277	49,126

Morgan Guaranty Trust Company
Risk Management Services
Peter Zangari (1-212) 648-8641
zangari_peter@jpmorgan.com

Finally, recall that the coverage cost of USD 870,994 is USD 8,289. Since this is a sunk cost, when computing the confidence bands it should be subtracted from the portfolio's expected return. Table 4 presents confidence bands which are adjusted for the both the portfolio's expected return and the option's cost.

Table 4

Confidence bands relative to expected portfolio return and option cost

values correspond to VaR + mean portfolio return - option cost (in USD)

Confidence Interval 90%	Normal	Cornish Fisher	Partial simulation	Full simulation
Lower 5%	-22,499	-20,036	-21,044	-12,297
Upper 5%	35,475	46,320	44,566	40,837

Conclusion

This article describes two alternative methodologies to estimate VaR on portfolios that include options when accounting for gamma risk. The Cornish-Fisher expansion provides an analytical approximation to the percentiles of a portfolio's true return distribution. In VaR estimation, these percentiles are used in place of their normal counterparts +/- 1.65. Based on a synthetic portfolio that consists of a French government bond (OAT) and a foreign exchange option, this approximation offers an improvement over normal VaR estimates. However, in practice, the Cornish-Fisher expression may yield inaccurate results because it is evaluated at sample estimates of skewness and kurtosis rather than at their true values. To address this issue, a partial simulation approach is suggested that requires the simulation of correlated multivariate normal random variates. These variates are then transformed into the portfolio's return distribution from which the appropriate percentiles are found. While partial simulation does not require the calculation of sample estimates, its main drawback is that the percentiles it produces are subject to simulation error. Ultimately, evaluating the performance of these adjustments is an empirical issue. This requires estimating VaR on portfolios of different size and composition. Those interested in such studies as well as the technical details of this paper including extensions to larger portfolios should contact the author using the number (e-mail address) listed above.